Re: Clustering Windows 2008

Source:
http://www.tech−archive.net/Archive/Windows/microsoft.public.windows.server.clustering/2008−07/msg00151.html

- *From*: "JJP" <anonymous@xxxxxxxxxxxxx>
- *Date*: Fri, 25 Jul 2008 09:28:40 −0500

Thanks Overstear for the reply...
I found out our EVA supports Windows 2008 but in the OS type on the EVA I had Microsoft Windows as the host type and I did not see the Microsoft Windows LH in the list. LH stands for LongHorn and there is an EVA controller code update that is supposed to change that to Microsoft Windows 2008 (I have not applied that yet). Anyway, once I selected "LH" for OS type, that error went away.

The first error was resolved by using the GUI to show hidden devices and disabling that Teredo device. For some reason doing it by command line would not work for us, but GUI did.

Now the cluster is up and I am just trying to figure it out....
Thanks again.
Jeff

"OVERSTEAR" <OVERSTEAR@xxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message news:42DE06FF−EF33−45D0−A363−C3AC42A167CA@xxxxxxxxxxxxxxx

> I can't help you with the first error but I can explain the second error.
> Windows 2008 uses Persistent SCSI reservations. The SCSI Reserve command and
> the SCSI Persistent Reserve command are specified by the SCSI standards.
> Servers can use these commands to prevent HBA ports in other servers from accessing the LUN.
> This prevents accidental data corruption that is caused when a server overwrites data on another server. This is why you can add your disks to both
> nodes before the cluster is configured. The Reserve and Persistent Reserve
> commands are often used by clustering software to control access to SAN Volume Controller virtual disks (VDisks).
> If a server is not shut down or removed from the server cluster in a controlled way, the server reserves and persistent reserves are

maintained.
This prevents other servers from accessing data that is no longer in use by
the server that holds the reservation. In this situation, you might want to
release the reservation and allow a new server to access the VDisk. If you
ever had a node in a cluster hang and not failover to other nodes in the
cluster this will help by letting a different node force access to the disk
and allow failover.
Now for the bad news your SAN must support SCSI reservations for Windows
2008 Failover Clustering. There are no exceptions. Have a new EMC array that
supports SCSI reservations and have built 3 Windows 2008 failover clusters.
Let me say W2008 Failover Clustering is leaps and bounds better than MSCS.
In my co−location we rent SAN disk from the provider and they don't support
persistent reservation and I am stuck using MSCS on window 2003. In case you
care my new EMC array will be installed in our co−location on Friday and I
will be migrating to Failover clustering A.S.A.P.

"JJP" wrote:

> Hi, we are trying to make a simple file share cluster with a couple of
> Windows 2008 64−bit servers. The Step By Step Guide from MS is not nearly
> as good and detailed for 2008 as it was for 2003.
>
> I presented 2 LUN's to both servers (which seems incorrect to begin with,
> but that is what the documents says) and both servers see the LUN's, MPIO is
> installed. So....
>
> I run the Validate a cluster tool.
>
> First error it gives:
> "Found duplicate IP address, blah blah blah...." This is the Teredo
> Tunneling Pseudo Interface. I found a previous post and ran the command
> "netsh interface teredo set state disabled", the server said "Ok" (whatever
> THAT means) and it still errors on the validation. I rebooted the potential
> nodes and it still errors on the validation.
>
> Second error it gives:
> "Validate storage supports SCSI−3 Persistent Reservation commands"
> "Failed to read drive layout of Cluster Disk 0 from node NodeNameB,

status
21"
"Cluster Disk 0 does not support Persistent Reservation"
"Cluster Disk 1 does not support Persisent Reservation"
So does that have to do with the fact that I presented both LUN's to both potential nodes without cluster installed? Usually I would not do that but
the step by step guide says 2 LUN's must be presented to the nodes. Then again it is trying to test the LUN's from both nodes so I don't know.

This is frustrating, especially because even if we get past these errors, all the other tests that didnt run because of these first errors will probably just create more errors.

Thanks in advance for any help.

.